

MOVABLE TAP FINITE IMPULSE RESPONSE FILTER

BACKGROUND OF THE INVENTION

5

Field Of The Invention

The present invention relates to a finite impulse response filter, and particularly to such a filter in which a delay in a portion thereof has an adjustable or selectable delay period, and to an echo canceller and an Ethernet transceiver including such an FIR filter.

Description Of The Related Art

15

Finite impulse response (FIR) filters are extremely versatile digital signal processors that are used to shape and otherwise to filter an input signal so as to obtain an output signal with desired characteristics. FIR filters may be used in such diverse fields as Ethernet transceivers, read circuits for disk drives, ghost cancellation in broadcast and cable TV transmission, channel equalization for communication in magnetic recording, echo cancellation, estimation/prediction for speech processing, adaptive noise cancellation, etc. For example, see U.S. Patent Nos. 5,535,150; 5,777,910; and 6,035,320, the contents of each of which are incorporated herein by reference. Reference is also made to the following publications: "An adaptive Multiple Echo Canceller for Slowly Time Varying Echo Paths," by Yip and Etter, IEEE Transactions on Communications, October 1990; "Digital Signal Processing", Alan V. Oppenheim, et al., pp. 155-163; "A 100MHz Output Rate Analog-to-Digital Interface for PRML Magnetic-Disk Read Channels in 1.2um CMOS", Gregory T. Uehara and Paul R. Gray, ISSCC94/Session 17/Disk-Drive

Electronics/ Paper FA 17.3, 1994 IEEE International
Solid-State Circuits Conference, pp. 280-281; "72Mb/S PRML
Disk-Drive Channel Chip with an Analog Sampled Data Signal
Processor", Richard G. Yamasaki, et al., ISSCC94/Session
5 17/Disk-Drive Electronics/Paper FA 17.2, 1994 IEEE
International Solid-State Circuits Conference, pp. 278,
279; "A Discrete-Time Analog Signal Processor for Disk
Read Channels", Ramon Gomez, et al., ISSCC 93/Session
13/Hard Disk and Tape Drives/Paper FA 13.1, 1993
10 ISSCC Slide Supplement, pp. 162, 163, 279, 280; and "A
50MHz 70 mW 8-Tap Adaptive Equalizer/Viterbi Sequence
Detector in 1.2 um CMOS", Gregory T. Uehara, et al. 1994
IEEE Custom Integrated Circuits Conference, pp. 51-54, the
contents of each being incorporated herein by reference.

15
Typically, an FIR filter is constructed in
multiple stages, with each stage including an input, a
multiplier for multiplication of the input signal by a
coefficient, and a summer for summing the multiplication
20 result with the output from an adjacent stage. The
coefficients are selected by the designer so as to achieve
the filtering and output characteristics desired in the
output signal. These coefficients (or filter tap weights)
are often varied, and can be determined from a least mean
25 square (LMS) algorithm based on gradient optimization.
The input signal is a discrete time sequence which may be
analog or digital, while the output is also a discrete
time sequence which is the convolution of the input
sequence and the filter impulse response, as determined by
30 the coefficients.

With such a construction, it can be shown
mathematically and experimentally that virtually any
linear system response can be modeled as an FIR response,

as long as sufficient stages are provided. Because of this feature, and the high stability of FIR filters, such filters have found widespread popularity and are used extensively.

5

One problem inherent in FIR filters is that each stage requires a finite area on an integrated circuit chip. Additional area is required for access to an external pin so as to supply the multiplication or
10 weighting coefficient for that stage. In some environments, the number of stages needed to provide desired output characteristics is large. For example, in Gigabit Ethernet applications it is preferred that every 8 meters of cable length be provided with 11 stages of FIR
15 filter. In order to cover cable lengths as long as 160 meters, 220 FIR stages should be provided. In such environments, chip area on the integrated circuit is largely monopolized by the FIR stages.

Moreover, each FIR stage requires a finite
20 amount of power and generates a corresponding amount of heat. Particularly where a large number of stages is needed, such power requirements become excessive and require significant mechanical adaptations to dissipate
25 the heat.

The inventors herein have recently recognized that in some environments, not all stages of an FIR contribute significantly to the output. Figure 1, for
30 example, is a waveform showing signal amplitude versus time in an Ethernet echo cancellation application, where time (on the horizontal axis) is expressed in delay units for an FIR filter. The waveform shown in Figure 1 represents an Ethernet transmission and its echo (or,
35 reflection). As seen in Figure 1, the waveform includes

the near end echo at region 1, followed by a relatively quiet period in region 2, a relatively negligible signal at region 3, and the far end echo at region 4. One use of an FIR filter in such an Ethernet environment is to cancel the echo so as to distinguish more clearly between incoming signals and simple reflections of transmitted signals. However, the relatively negligible signal at region 3 contributes very little to the overall output of the FIR filter. The reason for this is that, whatever value of coefficients are set at the stages corresponding to region 3, those coefficients will be multiplied by a value which is approximately zero. Thus, contributions of those signals to the FIR output will be negligible, especially compared to regions 1, 2 or 4.

The inventors have considered simplifying the selection of coefficients by setting the coefficients corresponding to region 3 to zero, which would result in simpler algorithms needed to select coefficients. However, even with zeroed coefficients, the stages corresponding to region 3 still exist on the integrated circuit chip, stealing valuable surface area and power, and generating unwanted heat.

SUMMARY OF THE INVENTION

It is an object of the present invention to address the foregoing, by providing an FIR filter in which the delay of one or more stages is selectable or adjustable with respect to the other stages.

By virtue of this arrangement, since the delay of one stage is adjustable, it is possible to "skip" areas of the input signal that are known to have negligible signal level relative to other areas of the input signal.

That is, the portion of the input signal may be delayed by a variable period before being injected into a predetermined block of FIR stages thus "skipping" over the irrelevant portions of the signal. Since the entire input signal does not need to be injected into FIR stages, fewer stages are required to filter the input signal. Moreover, since the "skipped" stages need not be fabricated on the chip, an FIR filter according to the invention ordinarily has reduced surface area, power requirements, heat generation, and taps for coefficients, relative to a prior art FIR filter in which one or more stages do not have an adjustable or selectable delay. For example, an Ethernet echo canceller that required 224 FIR stages at the prior art can now be constructed with 160 stages, i.e., 160 taps plus 64 virtual taps having a coefficient of approximately zero.

Thus, in one aspect, the invention comprises an FIR filter having multiple coefficient taps, each associated with an input signal in corresponding stages of delay from a corresponding delay element. At least one delay element has a period of delay that is selectable or adjustable independently of the period of delay for other delay elements. Preferably, the period of delay is selectable or adjustable through pin-out elements of the FIR filter. In the preferred embodiment, at power-up all delay elements in the FIR start with the same starting period of delay. Thereafter, the delay of one or more stages is adjustable with additional delay, meaning that the delay of that stage can be more, but no less than, the starting delay.

By virtue of the foregoing arrangement, in which one or more stages has a delay that is selectable or adjustable, an FIR filter according to the invention is

smaller, requires fewer pins, uses less power and generates less heat than conventional FIR filters. In particularly preferred embodiments, the FIR filter has many stages, such as 160 stages, with the delay for one or more stages being selectable to skip over relatively negligible areas of the input signal, thereby providing an FIR filter whose output performance is very close to that of an FIR filter with many more stages, such as 224 stages.

10

Another aspect of the invention concerns the method of selecting a delay in an FIR filter comprising multiple coefficient taps each associated with an input signal in corresponding stages of delay from a corresponding delay element in which at least one delay element has a period of delay that is selectable or adjustable independently of that of other delay elements.

15

According to the method, delayed components of a response signal are measured, so as to identify a sequence of components that are negligibly small compared to other sequences of components. The selectable period of delay is then set to a value calculated to prevent exposure of the identified sequence to the coefficient taps. In preferred embodiments, the response signal is monitored at the beginning of the transmission of data to determine the absolute maximum value of the far end echo tap location. Preferably, the center of the far end echo of region 4 is centered on the last section of the FIR filter taps, after the delayed portion.

20

25

30

Because the selectable period of delay is set to a value calculated to prevent exposure of the negligible sequence of components to the coefficient taps, the selected period of delay effectively "skips" negligible periods of the filtered input signal. Preferably, the

35

method for selecting a period of delay is iterative, meaning that a selectable period of delay is set based on a measurement of delayed components, performance of the FIR filter is checked based on the current period of delay
5 based on a further measurement of delayed components, and a next iterative period of delay is thereafter set based on the latest measurement of delayed components.

This brief summary has been provided so that the
10 nature of the invention may be understood quickly. A more complete understanding of the invention can be obtained by reference to the following detailed description of the preferred embodiments in connection with the attached drawings.

15

BRIEF DESCRIPTION OF THE DRAWINGS

Figure 1 is a view showing a channel response
20 waveform over copper cable in an Ethernet environment, including near end echo and far end echo due to reflection.

Figure 2 is a functional block diagram showing
25 an Ethernet transceiver including a transmit side and a receive side, and in which an echo canceller thereof includes an FIR filter according to the invention.

Figure 3 is a functional block diagram of the
30 echo canceller in Figure 2, showing an FIR filter according to the invention together with least mean square elements by which the coefficient for each stage is generated, and including an adjustable delay element.

Figure 4 is a functional block diagram of the 64-delay pipe shown in Figure 3.

Figures 5a and 5b are functional block diagrams showing the FIR filter of Figure 3.

Figure 6 is a functional block diagram showing the quantizer and downsampling blocks of the FIR filter of Figure 3.

Figure 7 is a flowchart depicting a method of determining how much delay should be provided to the input signal in accordance with the present invention.

DETAILED DESCRIPTION OF THE PRESENTLY PREFERRED EMBODIMENTS

The present invention will now be described with reference with to an echo canceller used in an Ethernet transceiver device. Preferably, the echo canceller is embodied in an Integrated Circuit disposed between a digital interface and an RJ45 analog jack. The Integrated Circuit may be installed inside a PC on the network interface card or the motherboard, or may be installed inside a network switch or router. However, other embodiments include applications in read circuits for disk drives, ghost cancellation in broadcast and cable TV transmission, channel equalization for communication in magnetic recording, echo cancellation, estimation/prediction for speech processing, adaptive noise cancellation, etc. All such embodiments are included within the scope of the appended claims.

While the present invention is described with respect to a digital FIR filter, is to be understood that the structure and functions described herein are equally applicable to an analog FIR. Moreover, while the invention will be described with respect to the functional elements of the FIR filter, the person of ordinary skill in the art will be able to embody such functions in discrete digital or analog circuitry, or as software executed by a general purpose process (CPU) or digital signal processor.

A functional block diagram of an Ethernet transceiver incorporating an FIR filter according to the present invention is depicted in Figure 2. Although only one channel is depicted therein, four parallel channels are typically used in Gigabit Ethernet applications. Only one channel is depicted and described herein for clarity.

A 125 MHz, 250Mbps digital input signal from a PC is PCS-encoded in a PCS encoder 2 and is then supplied to a D/A converter 4 for transmission to the Ethernet cable 6. The PCS-encoded signal is also supplied to a NEXT (Near End Transmitter) noise canceller 8 and to adaptive echo canceller 10. The operation of the echo canceller 10 will be described later herein with respect to Figure 3.

Signals from the Ethernet cable 6 are received at adder 14 and added with correction signals supplied from baseline wander correction block 12 (which corrects for DC offset). The added signals are then converted to digital signals in the A/D converter 16, as controlled by timing and phase-lock-loop block 18. The digital signals from A/D converter 16 are supplied to delay adjustment block 20, which synchronizes the signals in accordance

with the four parallel Ethernet channels. The delay-adjusted digital signals are then added with the echo-canceled signals and the NEXT-canceled signals in adder 22.

5

The added signals are supplied to a Feed Forward Equalizer filter 24 which filters the signal prior to Viterbi trellis decoding in decoder 26. After Viterbi decoding, the output signal is supplied to PCS decoder 28, after which the PCS-decoded signal is supplied to the PC.

The decoder 26 also supplies output signals to a plurality of adaptation blocks schematically depicted at 30 in Figure 2. As is known, such adaptation blocks carry out corrections for such conditions as temperature offset, connector mismatch, etc. The adaptation block 30 provides output to the baseline wander correction circuit 12, the timing and phase-lock-loop circuit 18, the echo canceller 10, and the NEXT canceller 8.

20

Each functional block depicted in Figure 2 includes a slave state controller (not shown) for controlling the operation and timing of the corresponding block. A PCS controller 32 controls the slave state controllers of all elements depicted in Figure 2, in a manner to be described below.

Figure 3 is a functional block diagram of the echo canceller 10 shown in Figure 2. In Figure 3, the PCS-encoded logic signal is provided to logic encoder 302 as a five level logic signal (e.g. -1, -0.5, 0, +0.5, +1).

The encoder 302 encodes the signal as 3 control bits, which correspond to the five logic levels of the PCS-encoded signal (e.g. -1=100; -0.5=101; 0=010; 0.5=001; 1=000). These control bits are supplied to a first

plurality or block of filter stages 304 (comprising taps 0 to 31 of the FIR filter), a second plurality or block of filter stages 306 (comprising taps 32 to 63), a third plurality or block of filter stages 308 (comprising taps 64 to 95), and a fourth plurality or block of filter stages 310 (comprising taps 96 to 127).

Filter blocks 304, 306, 308, and 310 typically have fixed delay periods between each of the taps for constant sampling of the early regions of the input signal where significant signal strength is present. Referring to Figure 1, large amplitudes are present in regions 1 and 2 of the input signal, and (according to the present embodiment) the blocks 304, 306, 308, and 310 receive these regions of the input signal to insure filtering of these significant portions of the signal. A more detailed description of the filter blocks will be provided later herein.

The logic-level-encoded signal from encoder 302 is also supplied to a 64-delay pipe (with 4 increment) 312. The delay pipe 312 is controlled by the echo controller's sequence control state machine 314 so that the portion of the input signal having the most significant echo noise is supplied to filter block 316 for noise cancellation. That is, the region 3 of the input signal is delayed appropriately in delay pipe 64 so that region number 3 is not subjected to echo cancellation (it is "skipped over") until portion 4 can be received and input into filter block 316. This way, not the entire input signal is FIR-filtered, and not as many taps are needed to effectively cancel the echo in the input signal.

The method by which the signal is selectively delayed will be described in more detail below.

The output of the logic level encoder 302 is also supplied to a quantizer 318 which encodes the three control bits into two logic bits for application to downsampling blocks 322 and 324 (to be described below).

5 For example, the quantizer 318 encodes 000 as 00; 001 as 00; 010 as 10; 101 as 01; and 100 as 01. The quantizer 318 thus performs a rounding function so that the encoded signal may be used to control the least mean squares (LMS) engines 0 through 6.

10

The LMS engines 4, 5, and 6 are designed to supply tap weighting coefficients to a single block of 32 FIR filter taps, and thus downsampling block 324 can use the same quantizer data for 32 cycles. In contrast, and
15 in accordance with the present invention, LMS engines 0, 1, 2, and 3 are designed to supply tap weighting coefficients to taps 0 to 31 of filter block 304, and downsampling block 322 controls each of these LMS engines in a time-cyclic fashion. This architecture allows more
20 precise filtering of the early regions of the input signal having significant signal strength. For example, at time t1, LMS engine 0 supplies a weighting coefficient to tap 0, LMS engine 1 supplies a weighting coefficient to tap 1, LMS engine 2 supplies a weighting coefficient to tap 2, and LMS engine 3 supplies a weighting coefficient to tap
25 3. At time t2, LMS engine 0 supplies a weighting coefficient to tap 1, LMS engine 1 supplies a weighting coefficient to tap 2, LMS engine 2 supplies a weighting coefficient to tap 6, and LMS engine 3 supplies a
30 weighting coefficient to tap 4. In this cyclic fashion, LMS engines 0-3 supply weighting coefficients to more precisely filter the region 1 of the input signal, in contrast to the less precise filtering of the region 2 of the input signal filtered by filter blocks 306, 308, and
35 310. The above is described in more detail in commonly

assigned U.S. Patent application Serial No. 09/465228,
filed December 19, 1999 and entitled, ``A Method and
Apparatus for Digital Near-End Echo / Near-End Crosstalk
Cancellation with Adaptive Correlation'', the contents of
5 which is incorporated herein by reference.

The quantizer 320 quantizes the output of the
delay pipe 312 and supplies it to the downsampling block
324 in a manner similar to that described above with
10 respect to quantizer 318. Downsampling block 326 then
controls LMS engine 7 which supplies weighting
coefficients to the taps 128 to 159 of the filter block
316 (which thus filters the adaptively delayed portion of
the input signal).

15 The manner by which the LMS engines generate the
tap coefficients will now be described. The LMS engines 0
to 7 input error signals from the FFE 24 or the Viterbi
decoder 26 of Figure 2. A memory 330 stores weighting
20 coefficients for each of taps 32-127. As the error signal
is received from the FFE 24 or the Viterbi decoder 26, the
appropriate coefficients are extracted from memory 330,
applied through the corresponding LMS engine, and provided
to the appropriate taps 32-127 in order to filter the
25 input signal to eliminate the echo noise in region 2 of
the input signal.

In a manner similar to that described above,
memory 332 stores coefficients for the taps 0-31 of the
30 filter block 304. The appropriate coefficients are
extracted from memory 332 and applied to the appropriate
LMS engines 0-3 together with the error signal, and the
appropriate coefficients are then supplied to the taps 0-
31 to appropriately filter the echo noise in region 1 of
35 the input signal. Similarly, the memory 334 stores

coefficients for the taps 128-159, which are selectively applied to the LMS engine 7 together with the error signal. The appropriate tap coefficients are then applied to filter block 316.

5

Figure 4 is a functional block diagram of the 64-delay element 312 of Figure 3. As can be seen, the 64 delay elements are grouped in sets of four delay elements 412, 414, 416, and 418. The logic level-encoded signal S is input to the delay pipe and may be delayed in increments of four by activation of control signals at gates 420, 422, and 424. The control signals are supplied by the sequence control state machine 314, and are varied in accordance with which portion of the input signal is to be skipped, as will be described below.

Figure 5a is a functional block diagram of the FIR filter showing how the variable delay D is supplied to an existing delay element 512 in order to variably adjust the input signal to skip the desired portion thereof. In Figure 5a, the logic level-encoded signal S is supplied, for example, to a first element 520 having a time delay t_1 . A tap coefficient C_0 is applied to a multiplier 505 in order to weight the first tap of the FIR filter. The weighted signal is then provided to a summer 515 where it is added to the outputs of the other stages (to be described below), and then output as signal S_o . The signal S is also supplied to the multiplier 518 for multiplication by coefficient C_1 , and addition with the other outputs at summer 514. Of course, any number of additional stages like 520 may be provided prior to the output, as required.

The input signal S is also supplied to delay element 512 having a variable delay D. The thus-delayed

09678728-100400

signal Svd is then provided to a series of sequential delay elements including delay element 506, which preferably also has a fixed delay time t_1 . The delayed signal Svd is also supplied to multiplier 516 for

5 multiplication by coefficient C_{n-2} and addition in summer 513, as shown. The output of delay element 506 $Svd+t_1$ is supplied to both another delay element 502 (having a t_1 delay) and a multiplier 510 where it is multiplied by coefficient C_{n-1} . The output of element 502 $Svd+t_1+t_1$ is

10 supplied to multiplier 504 where it is multiplied by coefficient C_n and then added, in adder 508, to the output from multiplier 510. In this manner, the series of weighted tap coefficients and corresponding input signals are processed through the FIR filter, in a manner known to

15 those of skill in the art.

The appropriate number of stages with corresponding delay elements are provided in order to properly filter the regions of the input signal having

20 significant signal strength, such as regions 1 and 2 in Figure 1. However, to skip those insignificant portions of the signal (such as region 3), the element 512 is provided with the variable delay D in accordance with control signal C_t supplied from the sequence control state

25 machine 314. According to the present invention, the variable delay D may be selected to skip any portion of the input signal which is not to be filtered. Preferably, a later portion of the input signal will be filtered since significant echo typically resides therein. Accordingly,

30 after element 512, any number of additional stages like elements 502 and 506 are provided, typically having the same fixed time delay t_1 . The number of additional stages after stage 512 may be varied to appropriately filter the echo regions of the input signal.

Figure 5b shows an alternative wherein the delay element 584 is provided to the undelayed portion of the input signal S to skip portions thereof. Like reference numerals represent like structure. In Figure 5b, the input signal S is supplied to both of multipliers 590 and 592 where it is respectively multiplied by coefficients C0 and C1. The delayed signal Svd output from element 584 is, after any number of intervening stages, supplied to both multipliers 510 and 504 where it is respectively multiplied by coefficients Cn-1 and Cn. The output of multiplier 504 is delayed in a delay element 502 having a t1 delay, and then supplied to adder 508 where it is added to the output from multiplier 510. The output of adder 508 is then supplied to a delay element 506 having a delay of t1, and the output of 506 is, in turn, provided (after any number of intermediate stages) to the adder 514 where it is added with the output of multiplier 590. The output of adder 514 is provided to a delay element 586 having a t1 delay. The output of the element 586 is added, in adder 588, to the output of multiplier 592, and the output of adder 588 is the output signal S0.

In a further alternative to the above arrangement, variable delays may be provided to more than one filter block. For example, filter block(s) 310 and/or 308 and/or 306 may also be supplied with variable delays so that any portions of the input signal may be skipped or filtered as the circuit designer requires. All such alternatives are included within the scope of the appended claims.

Figure 6 is a functional block diagram of the quantizer and downsampling circuits of Figure 3. The quantizer 318 receives the logical level-encoded signal S from the input of delay pipe 312. The output of quantizer

318 is provided to both the downsampling block 324 and a multiplexer 612. The multiplexer 612 outputs the quantizer signal to a one-cycle delay element 614, which supplies the down-sampled signal to LMS engine 3. In a similar manner, delay elements 616, 618, and 620 respectively provide down-sampled signals to LMS engines 2, 1, and 0, after the appropriate delay. The output of delay element 620 is also returned to the multiplexer 612, as shown.

10

The output of downsampling block 324 is provided to the LMS engines 6, 5, and 4, as was described above with reference to Figure 3. Also, the output of the delay pipe 312 is supplied to the quantizer 320 which supplies the downsampling block 326 and LMS engine 7, as shown.

In operation, those portions of the input signal which may be skipped by the FIR filter must first be determined. Preferably, this is done by injecting a test signal into the Ethernet cable and then receiving the return signal, such as the waveform depicted in Figure 1. However, the procedure for determining the insignificant portions of the input signal may be performed at any convenient time, such as when the Ethernet is first powered on, after any Ethernet device has been plugged into the network or unplugged from the network, during any lull in Ethernet communications, on a periodic basis, or continually. The signal used to determine the delay may also be any appropriate signal such as a test signal, a series of test signals, or by using actual Ethernet communication signals on-the-fly.

The method of determining how much delay to be supplied to the input signal in accordance with the embodiment of Figure 3 will now be described with respect

to the flow chart of Figure 7. This process is preferably carried out within the sequence control state machine 314, although any convenient processor and memory may be used. In Figure 7, when the Ethernet is first powered-up, data
5 starts to be supplied to the Ethernet cable 6 at step S1. At step S2, the return signal is received and then filtered in the FIR filter using blocks 304, 306, 308, 310, and 316 contiguously so as to filter a continuous portion of the return signal. At step S4, it is
10 determined which tap of taps 128-159 has received the maximum return signal strength. This tap is labeled tapmaxd. At step S5, tapmaxd is compared with the stored tapmaxs, and the tap having the maximum signal strength is then stored as the new tapmaxs. Of course, for the first
15 determination, the initial tapmaxd will be stored as tapmaxs. In order to avoid storing unexpectedly large signal strength caused by noise, multiple looping for comparison is preferably employed. For example, if 32 taps are compared and tap 7 is identified as tapmaxs, the
20 comparison will be repeated multiple times. Every comparison, tap 7 will be replaced with tapmaxs even though the tapmaxs is larger than tap 7, in order to avoid a lock up error.

25 At step S6, it is determined whether the end of the return signal has been reached. If the end of the return signal has not been reached, the process proceeds to step S7 where a 32 tap delay is applied to skip a portion of the return signal. Of course, any amount of
30 tap delay (1 tap, 4 taps, 8 taps, 16 taps, 64 taps, etc.) may be used in any combination by the circuit designer to flexibly configure the FIR filter. The process then returns to step S4 to determine which tap of the newly-filtered signals has the maximum signal strength. Again,

the determined tapmaxd is compared with the stored tapmaxs, and the maximum value is stored as the new tapmaxs in step S5.

5 One algorithm for performing steps S4, S5, S6, and S8 of Fig. 7 is as follows:

Let n = the number of stages in the FIR filter.

Let tap[i] = the ith stage of the FIR filter.

10 Let {tap[i]} = the coefficient value of the ith stage of the FIR filter.

Let Maxcoeff = the absolute value of the maximum coefficient value in the FIR filter.

15 Let m = the index of which tap coefficient is written into Maxcoeff.

At time = 0,

Maxcoeff ← {tap[0]}

m ← 0

20

At time = i, (where i > 0, i.e., 1, 2, 3, 4,...)

if (en_search) //where en_search enables the search for Maxcoeff
begin

if (Maxcoeff < {tap[i]} or m = i)

25

begin

Maxcoeff ← {tap[i]}

m ← I

end

else

30

begin

Maxcoeff ← Maxcoeff

m ← m

end

end

```
else
    begin
        Maxcoeff ← Maxcoeff
        m ← m
5    end.
```

00400T" B2782960

10 In this iterative manner, the last filter block 316 is successively moved across the later portions of the return signal identifying which portion(s) of the return signal have the maximum signal strength. When the filter block 316 has reached the end of the return signal, step S8 is performed wherein the stored tapmaxs is set as the

15 center tap of the filter block 316. Now, the filter block 316 will be applied to the center of the later portion of the return signal having the most significant signal strength. The required delay may be determined

20 algorithmically or from accessing an entry from a lookup table. The delay required to so-position filter block 316 is then stored in the memory of sequence control state machine 314 so that all Ethernet signals received from the Ethernet cable 6 may be FIR-filtered in accordance with

25 the thus-configured filter blocks to skip those portions of the signal having insignificant signal strength, while filtering the remaining signal. In such a manner, Ethernet signals typically requiring more than 220 taps for proper FIR filtration can be adequately filtered with an FIR filter having only 160 taps.

30

Thus, what has been described is method and apparatus for controlling an FIR filter so as to delay the input signal to skip over portions of that signal having insignificant signal strength. This allows the FIR filter

to have fewer taps, consuming less power and less space on the Integrated Circuit.

5 The individual components shown in outline or designated by blocks in the attached Drawings are all well-known in the FIR filtering arts, and their specific construction and operation are not critical to the operation or best mode for carrying out the invention.

10 While the present invention has been described with respect to what is presently considered to be the preferred embodiments, it is to be understood that the invention is not limited to the disclosed embodiments. To the contrary, the invention is intended to cover various
15 modifications and equivalent arrangements included within the spirit and scope of the appended claims. The scope of the following claims is to be accorded the broadest interpretation so as to encompass all such modifications and equivalent structures and functions.

20

25

30

35